graphic signature scheme to the document, and transmitting the resulting document, now a certificate of the temporal existence of the original document, back to the author where it is held for later use in required proof of such existence.

To ensure against interception of confidential document information during transmission, and to reduce the digital bandwidth required for transmission of the entire document, the author may optionally convert the digital document string to a unique number having vastly reduced digital size by means of a deterministic function which may, for example, be any one of a number of algorithms known in the art as "oneway hash functions". Such an application of hash functions has been described, among others, by Damgard in his discussions on the improvement of security in document signing techniques ("Collision-Free Hash Functions and Public Key Signature Schemes", *Advances in Cryptology—Eurocrypt* '87, Springer-Verlag, LNCS, 1988, Vol. 304, pp. 203–217). In practice of the present invention, however, the "one-way" characteristic typical of a hashing algorithm serves an additional purpose; that is, to provide assurance that the document cannot be revised subsequent to the time the TSA applies its time stamp.

A hashing function provides just such assurance, since at the time a document is hashed there is created a representative "fingerprint" of its original content from which it is virtually impossible to recover that document. Therefore, the time-stamped document is not susceptible to revision by any adversary of the author. Nor is the author able to apply an issued time-stamp certificate to a revised form of the document, since any change in the original content, even to the extent of a single word or a single bit of digital data, results in a different document that would hash to a completely different fingerprint number. Although the original document can thus not be recovered from the hashed document, a purported original document can nonetheless be proven by the fact that a true copy of the original document will always hash, assuming use of the same hashing algorithm, to the original number contained in the certificate.

Any available deterministic function, e.g. a one-way hash function such as that described by Rivest ("The MD4 Message Digest Algorithm", *Advances in Cryptology—Crypto,* '90, Springer-Verlag, LNCS, to appear), may be used in the present procedure. In the practice of the invention, such a hashing operation would normally be employed by the author to obtain the noted benefit of transmission security, although it might be effected by the TSA if the document were received in plaintext form. In whatever such manner the document content and incorporated time data are fixed against revision, there remains the further step, in order to promote the credibility of the system, of certifying to the members of an as yet unidentified universe that the receipt was in fact prepared by the TSA, rather than by the author, and that the time indication is correct, i.e. that it has not, for instance, been fraudulently stated by the TSA in collusion with the author.

To satisfy the former concern, the TSA uses a verifiable signature scheme, of a type such as the public key method earlier noted, to certify the time-stamp prior to its transmittal to the author. Confirmation of the signature at a later time, such as by decryption with the TSA's public key, proves to the author and to the universe at large that the certificate originated with the TSA. Proof of the veracity of the time-stamp itself, however, relies upon a following additional aspect of the invention.

One embodiment of this segment of the process, as generally depicted in FIG. 2, draws upon the relatively continuous flow of documents from the universe of authors through the facilities of the TSA. For each given processed document D_k, the TSA generates a time-stamp receipt which includes, for example, a sequential receipt number, r_k, the identity of the author, A_k, by ID number ID_k, or the like, the hash, H_k, of the document, and the current time, t_k. In addition, the TSA includes the receipt data of the immediately preceding processed document, D_{k-1}, of author, A_{k-1}, thereby bounding the timestamp of document, D_k, in the "past" direction by the independently established earlier receipt time, t_{k-1}. Likewise, the receipt data of the next received document, D_{k+1}, are included to bound the time-stamp of document, D_k, in the "future" direction. The composite receipt, now containing the time data of the three, or more if desired, sequential time-stamp receipts, or identifying segments thereof, is then certified with the cryptographic TSA signature and transmitted to the author, A_k. In like manner, a certificate containing identifiable representations of D_k and D_{k+2} would be transmitted to author, A_{k+1}. Thus, each of the time-stamp certificates issued by the TSA is fixed in the continuum of time and none can be falsely prepared by the TSA, since a comparison of a number of relevant distributed certificates would reveal the discrepancy in their sequence. So effective is such a sequential fixing of a document in the time stream that the TSA signature could be superfluous in actual practice.

A second embodiment of the invention, shown generally in FIG. 3, distributes the time-stamping task randomly among a broad universe, for example the multiplicity of authors utilizing the time-stamping process. A TSA could still be employed for administrative purposes or the requesting author could communicate directly with the selected time-stamping author/agents. In either event, the above-mentioned need for assurance that a time-stamp has not been applied to a document through collusion between the author and the stamping agency is met in the combination of the reasonable premise that at least some portion of the agency universe is incorruptible or would otherwise pose a threat of exposure to an author attempting falsification, and the fact that the time-stamping agencies for a given document are selected from the universe entirely at random. The resulting lack of a capability on the part of the author to select a prospective collusive agent of the author's own choosing substantially removes the feasibility of intentional time falsification.

The selection of the individual universe members who will act as the predetermined number of agents is accomplished by means of a pseudorandom generator of the type discussed by Impagliazzo, Levin, and Luby ("Pseudorandom Generation From One-Way Functions", *Proc. 21st STOC,* pp. 12–24, ACM, 1989) for which the initial seed is a deterministic function, such as a hash, of the document being time-stamped. Given as a seed input the document hash or other such function, the implemented pseudorandom generator will output a series of agency IDs. This agency selection is for all practical purposes unpredictable and random.

Once the agents are selected, the time-stamping proceeds as previously indicated with the exception that each agent individually adds the current time data to the